arXiv:1510.02324v1 [math.CO] 8 Oct 2015

On the structure of (banner, odd hole)-free graphs*

Chính T. Hoàng[1]

[1] Department of Physics and Computer Science, Wilfrid Laurier University, Waterloo, Ontario, Canada, N2L 3C5

choang@wlu.ca

Abstract

A hole is a chordless cycle with at least four vertices. A hole is odd if it has an odd number of vertices. A banner is a graph which consists of a hole on four vertices and a single vertex with precisely one neighbor on the hole. We prove that a (banner, odd hole)-free graph is either perfect, or does not contain a stable set on three vertices, or contains a homogeneous set. Using this structure result, we design a polynomial-time algorithm for recognizing (banner, odd hole)-free graphs. We also design polynomial-time algorithms to find, for such a graph, a minimum coloring and largest stable set.

1 Introduction

A *hole* is a chordless cycle with at least four vertices. An *antihole* is the complement of a hole. A hole is *odd* if it has an odd number of vertices. A graph is odd-hole-free if it does not contain, as an induced subgraph, an odd hole. Odd-hole-free graphs are studied in connection with perfect graphs (definitions not given here will be given later.) It follows from a result of Tuza, Král, Kratochvíl, and Woeginger ([18]) that it is NP-hard to color an odd-hole-free graph. In contrast, there is a polynomial-time algorithm (Grötschel, Lovász and Schrijver [8]) to find a minimum coloring of a graph with no odd holes and no odd antiholes.

Chudnovsky, Cornuéjols, Liu, Seymour, and Vuskovic ([2]) designed a polynomial-time algorithm for finding an odd hole or odd antihole, if one exists, in a graph. However, the complexity of recognizing odd-hole-free graphs is unknown. Conforti, Cornuéjols, Kapoor, and Vušković ([5]) found a decomposition of odd-hole-free graphs by "double star-cutsets" and "2-joins" into certain "basic graphs". But this decomposition does not seem to help in designing a polynomial-time recognition algorithm for odd-hole-free graphs.

*Research support by Natural Sciences and Engineering Research Council of Canada.

A *banner* is the graph which consists of a hole on four vertices and a single vertex with precisely one neighbor on the hole. Banner-free graphs generalize the well studied class of claw-free graphs. In this paper, we study (banner, odd hole)-free graphs. We will show that if a (banner, odd hole)-free graph is not perfect, then either it contains a homogeneous set or its stability number is at most two. We will use this structural result to design polynomial-time algorithms for (i) recognizing a (banner, odd hole)-free graph, (ii) finding an optimal coloring of a (banner, odd hole)-free graph, and (iii) finding a largest stable set of a (banner, odd hole)-free graphs. In contrast, results in the literature show that, for a (banner, odd hole)-free graph, finding a largest stable set and finding a minimum cover by cliques are both NP-hard problems. We will discuss these facts in Section 6. We note the algorithm of Gerber, Hertz, and Lozin ([7]) that finds in polynomial time the largest stable set of a (banner, P_8)-free graph. Also, the result of Tuza et al. ([18]) shows that it is NP-hard to find a minimum coloring of an odd-hole-free graph.

The notion of "robust" algorithms were introduced by Raghavan and Spinrad ([16]). A *robust* algorithm for problem P on domain C must solve P correctly for every input in C. For input not in C, the algorithm may produce correct output for problem P, or answer that the input is not in C. Our optimization algorithms are robust.

A graph G with at least one edge is *k-divisible* if the vertex-set of each of its induced subgraph H with at least one edge can be partitioned into k sets, none of which contains a largest clique of H. It is easy to see that the chromatic number of a k-divisible graph is at most 2^{k-1}. It was conjectured by Hoàng and McDiarmid ([9]) that every odd-hole-free graph is 2-divisible. They proved the conjecture for claw-free graphs ([9]). We will prove the conjecture for banner-free graphs.

In Section 2, we give the definitions used in this paper and discuss background results that are used by our algorithms. In Section 3, we prove our theorems on the structure of (banner, odd hole)-free graphs. In Section 4, we prove that (banner, odd hole)-free graphs are 2-divisible. In Section 5, we give two polynomial-time algorithms for recognizing a (banner, odd hole)-free graph. In Section 6, we give polynomial-time algorithms for finding a minimum coloring and a largest stable set of a (banner, odd hole)-free graph.

2 Definitions and background

2.1 Definitions

Recall the definitions of the claw and banner (see Figure 1). Let K_t denote the clique on t vertices. Let C_k (respectively, P_k) denote the cordless cycle (respectively, path) on k vertices. The *girth* of a graph is the length of its shortest cycle. A *hole* is the graph C_k with $k \geq 4$. A hole is *odd* if it has an

(a) claw (b) banner

Figure 1: The claw and the banner

odd number of vertices. An *antihole* is the complement of a hole. A clique on three vertices is called a *triangle*. For a given graph H, it is customary to let *co-H* denote the complement of H. Thus, a co-triangle is the complement of the triangle. Let L be a list (set) of graphs. A graph G is L-free if G does not contain an induced subgraph isomorphic to any graph in L. In particular, a graph is *(banner, odd hole)-free* if it does not contain an induced subgraph isomorphic to a banner or an odd hole. Let G_1, G_2 be two vertex-disjoint graphs. The union of G_1 and G_2 is denoted by $G_1 + G_2$. If k is an integer, then kG_1 denotes the union of k disjoint copies of G_1. For a graph G and a set X of vertices of $V(G)$, $G[X]$ denotes the subgraph of G induced by X.

Let G be a graph. Then $\chi(G)$ denotes the chromatic number of G, and $\omega(G)$ denotes, the *clique number* of G, that is, the number of vertices in a largest clique of G. A set H of a graph G is *homogeneous* if $2 \leq |H| < |V(G)|$ and every vertex in $G - H$ is either adjacent to all vertices of H, or non-adjacent to any vertex of H. A graph is *prime* if it contains no homogeneous set. In algorithm analysis, it is customary to let n, respectively, m, denote the number of vertices, respectively, edges, of the input graph.

A graph G is *perfect* if $\chi(H) = \omega(H)$ for every induced subgraph H of G. Two important results are known about perfect graphs. The Perfect Graph Theorem, proved by Lovász ([12]), states that a graph is perfect if and only if its complement is. The Strong Perfect Graph Theorem, proved by Chudnovsky, Robertson, Seymour, and Thomas ([3]), states that a graph is perfect if and only if it is odd-hole-free and odd-antihole-free. Both of the above results were long standing open problems proposed by Berge ([1]). Chudnovsky et al. ([2]) designed an $O(n^9)$ algorithm for finding an odd hole or odd antihole in a graph, that is, for recognizing a perfect graph. Grötschel et al. ([8]) designed a polynomial-time and robust algorithm for finding a largest clique and a minimum coloring of a perfect graph. Corollary 1 in Král et al. ([18]) shows that it is NP-hard to compute the chromatic number of a $(2K_2, K_2+2K_1, 4K_1, C_5)$-free graph. It was pointed out to the author by K. Cameron that this result implies that it is NP-hard to color an odd-hole-free graph.

2.2 Modular Decomposition

Our recognition algorithms use the well-studied modular decomposition which we now discuss. Let G be a graph. A *module* is a set M of vertices such that every vertex in $G - M$ is adjacent to either all vertices, or no vertices, of M. Trivially, $\{x\}$ for any $x \in V(G)$, and $V(G)$ are modules. A module M is *non-trivial* if $2 \leq |M| < |V(G)|$, that is, a non-trivial module is a homogeneous set. (Here, we are bound to use two different names for essentially the same structure. In the study of graph coloring, "homogeneous set" is often used, whereas in graph algorithm design, "module" is used.)

Two sets X and Y *overlap* if $X - Y, Y - X$, and $Y \cap X$ are all non-empty. $X \Delta Y$ denotes the symmetric difference of X and Y, that is, $(X - Y) \cup (Y - X)$. It is easy to see that if X and Y are overlapping modules then the following sets are also modules: $X - Y, Y - X, X \cup Y, X \cap Y$, and $X \Delta Y$. A module is *strong* if it does not overlap another module. A proper submodule M of a graph Q is *maximal* if there does not exist module M' of Q such that $M \subset M' \subset Q$. A maximal module is a strong module.

Modular decomposition refers to the process of partitioning the vertices of a graph G into its maximal modules $\mathcal{P} = (M_1, M_2, \ldots, M_t)$. Each M_i, if it is not prime, is recursively decomposed. The procedure stops when every module has a single vertex. To a modular partition $\mathcal{P}$, we associate a *quotient* graph $G_{\mathcal{P}}$ whose vertices are the modules defined in $\mathcal{P}$; two vertices v_i and v_j of $G_{\mathcal{P}}$ are adjacent if and only if the corresponding modules M_i and M_j are adjacent in G. For the graph in Figure 2, the modular decomposition is $\mathcal{P} = \{\{1\}, \{2, 3\}, \{4, 5\}, \{6, 7, 8, 9, 10\}\}$. The quotient graph $G_{\mathcal{P}}$ is the P_4. For each module M in $\mathcal{P}$, if M is not prime, then M is recursively decomposed into maximal modules. For example, the module $\{6, 7, 8, 9, 10\}$ is decomposed into $\{\{6\}, \{7\}, \{8\}, \{9, 10\}\}$ The result of the modular decomposition can be represented as a tree.

A module which induces a disconnected subgraph in the graph is a *parallel module*. A module which induces a disconnected subgraph in the complement of the graph is a *series module*. A module which induces a connected subgraph in the graph as well as in the complement of the graph is a *neighborhood module*.

If the current set Q of vertices induces a disconnected subgraph, Q is decomposed into its components. A node labeled P (for parallel) is introduced, each component of Q is decomposed recursively, and the roots of the resulting subtrees are made children of the P node. If the complement of the subgraph induced by current set Q is disconnected, Q is decomposed into the components of the complement. A node labeled S (for series) is introduced, each component of the complement of Q is decomposed recursively, and the roots of the resulting subtrees are made children of the S node. Finally, if the subgraph induced by the current set Q of vertices and its complement are connected, then Q is decomposed into its maximal proper submodules; it is known (Gallai [6]) that in this case, each vertex of Q belongs to a unique maximal proper submodule

of Q. A node labeled N (for neighborhood) is introduced, each maximal proper submodule of Q is decomposed recursively, and the roots of the resulting subtrees are made children of the N node. A graph and its modular decomposition tree are shown in Figure 2.

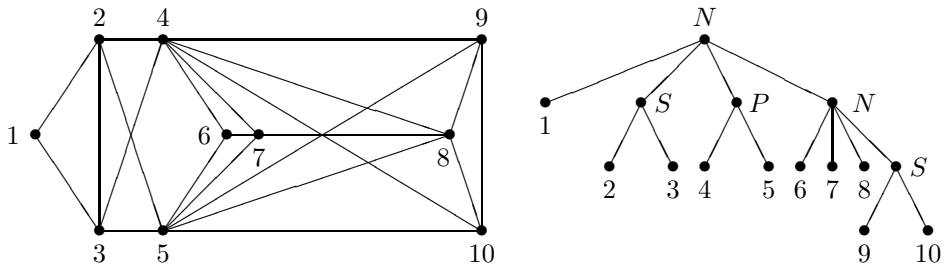

Figure 2: A graph and its modular decomposition tree

Theorem 2.1 (McConnell and Spinrad [13]) *The modular decomposition tree of a graph is unique and it can be constructed in $O(m+n)$ time.* □

To analyze our algorithms, it will be more convenient to consider the modular decomposition as a "binary" decomposition. Let H be a homogeneous set of a graph G. The decomposition produces two graphs: H and $G[(V(G)-H)\cup\{h\}]$ for a vertex $h \in H$. Note that $G[(V(G)-H)\cup\{h_1\}]$ is isomorphic to $G[(V(G)-H)\cup\{h_2\}]$ for any two vertices h_1, h_2 in H. If either of the graphs H and $G[(V(G)-H)\cup\{h\}]$ is not prime, then it is recursively decomposed. This process is summarized by the algorithm 1. Note that the homogeneous sets considered by the algorithm are not necessarily maximal. We will reproduce an argument used by Hoàng and Reed ([10]) to show that the number of prime graphs produced by Algorithm 1 is at most $n-1$ for an input graph with at least two vertices. Let $p(G)$ be the number of prime graphs produced by Algorithm 1 on a graph G with $|G|$ vertices. We are going to prove, for a graph G with at least two vertices, $p(G) \leq |G|-1$ by induction on the number of vertices of G. If G is prime, then $p(G)=1$. Let H be a homogeneous set of G that is used by Algorithm 1 in the decomposition step. The algorithm will decompose H and and $L = G[(V(G)-H\cup\{h\}\}$ for a vertex h of H. By the induction hypothesis, we have $p(G) = p(H)+p(F) \leq (|H|-1)+(|F|-1) = (|H|+|F|)-2$. Since $|H|+|F| = |G|+1$, we have $p(G) \leq |G|-1$.

3 The structure of (banner, odd hole)-free graphs

Chvátal and Sbihi ([4]) designed a polynomial-time algorithm for recognizing claw-free perfect graphs. In the process, they found a proof of the "Ben Rebea's

Algorithm 1 HOMOGENEOUS-SET-DECOMPOSITION

input: graph G
output: a set L containing prime induced subgraphs of G produced by modular decomposition.
initialization: $L \leftarrow \emptyset$, $W \leftarrow \{G\}$ (W contains the graphs to be decomposed)

if $W = \emptyset$ **then**
  stop
end if
repeat
  take one graph F from the set W and remove F from W
  if F is prime **then**
    add F to L
  else
    find a homogeneous set H of F
    add to W the two graphs: H and $F[(V(F) - H \cup \{h\}\}$ for an arbitrary vertex $h \in H$
  end if
until W is empty

Lemma" below.

Lemma 3.1 (Ben Rebea's Lemma [4]) *Let G be a connected claw-free graph with $\alpha(G) \geq 3$. If G contains an odd antihole, then G contains a C_5.*

It is this Lemma that inspires our two Lemmas 3.2 and 3.3 and Theorem 3.4 below.

Lemma 3.2 *Let G be a (banner, C_5)-free graph containing an odd antihole A such that*

(i) *$\alpha(G) \geq 3$, and*

(ii) *no co-triangle of G contains two vertices of A.*

Then the complement $\overline{G}$ of G contains a triangle-free component O that contains all of A.

Proof: Enumerate the vertices of A as $v_1, v_2, \ldots, v_k$ such that $v_i v_{i+1}$ is a non-edge of G with the subscripts taken module k. We will first establish the following claim.

No vertex in A belongs to a co-triangle. (1)

Suppose some vertex v_i of A forms a co-triangle with some two vertices u_1, u_2 in $G - A$. By (ii), we may assume v_{i+1} is adjacent to both u_1, u_2. Similarly,

v_{i-1} is adjacent to both u_1, u_2. (2)

We will show

Each of the vertices in $\{v_{i-2}, v_{i+2}\}$ is non-adjacent to at least one vertex in $\{u_1, u_2\}$. (3)

If v_{i+2} is adjacent to both u_1, u_2, then the graph $G[\{v_i, v_{i+1}, v_{i+2}, u_1, u_2\}]$ induces a banner. So, v_{i+2} is non-adjacent to at least one vertex in the set $\{u_1, u_2\}$. By symmetry, (3) holds. Next, we will show

Each of the vertices in $\{v_{i-2}, v_{i+2}\}$ is non-adjacent to all vertices in $\{u_1, u_2\}$. (4)

Suppose v_{i+2} is adjacent to u_2 and non-adjacent to u_1. We have the following implications.

$u_2v_{i+3} \in E(G)$, for otherwise $G[\{v_i, v_{i+1}, v_{i+2}, v_{i+3}, u_2\}]$ induces a C_5,

$u_1v_{i+3} \notin E(G)$, for otherwise $G[\{v_i, v_{i+2}, v_{i+3}, u_2, u_1\}]$ induces a banner,

$u_1v_{i-2} \in E(G)$, for otherwise $G[\{v_{i-2}, v_{i-1}, v_{i+2}, v_{i+3}, u_1\}]$ induces a banner,

$u_2v_{i-2} \notin E(G)$ by (3).

Now, $G[\{v_{i-2}, v_{i+2}, v_{i+3}, u_1, u_2\}]$ induces a banner. We have established (4).

It follows from (4) that, for all j, v_{i+2j} is non-adjacent to both u_1, u_2 with the subscripts taken modulo k. Since A is odd, this implies v_{i-1} is not adjacent to both u_1, u_2, a contradiction to (2). We have established (1).

Next, we will show

For any co-triangle C of $G - A$, some $v_i \in A$ is adjacent to all vertices of C. (5)

Suppose that (5) is false for some co-triangle C of $G - A$ with vertices t_1, t_2, t_3. By (1), any vertex v_i of A is adjacent to exactly two vertices of C.

For any subscript i, the vertices v_i, v_{i+1} cannot have the same neighbourhood in C, (6)

for otherwise v_i, v_{i+1} form a co-triangle with their unique non-neighbour in C, a contradiction to (ii).

Consider the vertices v_1, v_2, v_3. Suppose all of them have different neighbourhoods in C. We may assume without loss of generality that $N(v_1) \cap C = \{t_1, t_2\}, N(v_2) \cap C = \{t_1, t_3\}, N(v_3) \cap C = \{t_2, t_3\}$. Now $G[\{v_1, t_1, v_2, t_3, v_3\}]$ induces a C_5. Thus, the vertices v_1, v_2, v_3 all have distinct neighborhoods in C.

It follows from (6) that, for any i, v_i and v_{i+2} have the same neighbourhood in C. Since A has an odd number of vertices, all vertices in A have the same neighbourhood in C. But this is a contradiction to (6). We have established (5).

Let T be the set of vertices belonging to a co-triangle of G. By (1), we have $A \cap T = \emptyset$. Write $R = V(G) - (T \cup A)$. We will establish the following.

Let C be a co-triangle of G, and let v, x be two vertices in $G - T$. Suppose v is adjacent to all vertices of C. If x is not adjacent to v, then x is adjacent to all vertices of C. (7)

Let the vertices of the co-triangle C be t_1, t_2, t_3. Suppose x is non-adjacent to v and t_3. Since $x \notin T$, it is adjacent to t_1, t_2. But now $G[\{v, x, t_1, t_2, t_3\}]$ induces a banner in G. We have established (7).

Note (5) and (7) imply each vertex of A is adjacent to all vertices of T. Furthermore, every vertex in $G - T$ that has a non-neighbour in A is adjacent to all vertices of T. Now, consider the complement $\overline{G}$ of G and the component O of $\overline{G}$ that contains all of A. We claim that O contains no vertex of T, in particular, $\overline{G}$ is disconnected and O is triangle-free. Suppose in $\overline{G}$, there is a path with one endpoint in A and the other endpoint in T. Take such a shortest path P and enumerate the vertices of the path as $p_1, p_2, \ldots, p_j$ with $p_1 \in A, p_j \in T$ and $p_\ell \in V(G) - (T \cup A)$ for $\ell = 2, \ldots, j-1$. By (7), p_ℓ is adjacent, in G, to all of T for $\ell = 2, \ldots, j-1$; but this contradict the existence of the non-edge $p_{j-1}p_j$ (in G) of P. □

Lemma 3.3 *Let G be a (banner, C_5)-free graph with an odd antihole A. Suppose some two vertices of A belong to a co-triangle. Then there is a homogeneous set H of G such that H contains A and no co-triangle of $G[H]$ contains two vertices of A.*

Proof: By induction on the number of vertices. Let G and A be defined as in the Lemma. Let Q be the set of vertices q such that q forms a co-triangle with some two vertices of A. The set Q is not empty by the hypothesis of the Lemma. We may also assume G is connected, for otherwise, by the induction hypothesis, the component of G containing A contains the desired homogeneous set H. Enumerate the vertices of A as $v_1, v_2, \ldots, v_k$ such that $v_i v_{i+1}$ is a non-edge of G with the subscripts taken module k. We will first establish the claim below.

No vertex in Q has a neighbor in A. (8)

Suppose some vertex $u \in Q$ has some neighbour in A. Since u is non-adjacent to some two consecutive vertices of the antihole, we may assume $uv_1, uv_2 \notin E(G)$, and $uv_3 \in E(G)$. We have the following implications:

$uv_4 \notin E(G)$, for otherwise $G[\{v_1, v_2, v_3, v_4, u\}]$ induces a banner,

$uv_5 \notin E(G)$, for otherwise $G[\{v_1, v_2, v_4, v_5, u\}]$ induces a banner,
$uv_6 \in E(G)$, for otherwise $G[\{v_2, v_3, v_5, v_6, u\}]$ induces a banner.

But now $G[\{v_1, v_2, v_5, v_6, u\}]$ induces a banner. We have established (8).

Let R be the set of vertices in $V(G) - (A \cup Q)$ that have some neighbors in Q, and let S be the set of vertices in $V(G) - (A \cup Q)$ that have no neighbors in Q. Since G is connected, R is non-empty. Note that $V(G) = A \cup Q \cup R \cup S$. We will show that

Each vertex in $R \cup S$ is adjacent to some two consecutive vertices of A. (9)

Consider a vertex $x \in R \cup S$. Suppose x is not adjacent to some vertex v_i of A. Since $x \notin Q$, it must be adjacent to v_{i-1}, v_{i+1}. Thus, for any two consecutive vertices of A, the vertex x must be adjacent to at least one of them. Since A has an odd number of vertices, x must be adjacent to some two consecutive vertices of A. We have established (9). Next, we show that

Each vertex in R is adjacent to all vertices of A. (10)

Consider a vertex x in R. It is adjacent to some vertex u in Q. By (9), x is adjacent to some two consecutive vertices of A, say, v_1, v_2. For any $j \in \{4, 5, \ldots, k-1\}$, x is adjacent to v_j, for otherwise $G[\{v_j, v_1, x, v_2, u\}]$ induces a banner. Now, x is adjacent to v_3, for otherwise $G[\{v_{k-1}, v_{k-2}, v_3, x, u\}]$ induces a banner. Similarly, x is adjacent to v_k. We have establish (10).

Each vertex in R is adjacent to all vertices of S. (11)

Suppose there are vertices $r \in R, s \in S$ with $rs \notin E(G)$. By (9), s is adjacent to some two consecutive vertices v_i, v_{i+1} of A. By (10), r is adjacent to v_i, v_{i+1}. But now $G[\{v_i, v_{i+1}, r, s, u\}]$ induces a banner for some neighbor u in Q of r.

Now, the set $A \cup S$ is the desired homogeneous set of G. □

Theorem 3.4 *Let G be a (banner, odd hole)-free graph. Then one of the following holds for G.*

(i) *G is perfect.*

(ii) *$\alpha(G) \leq 2$.*

(iii) *Every odd antihole A of G belongs to a homogeneous set H of G such that $G[H]$ is co-triangle-free.*

Proof: Let G be a (banner, odd hole)-free graph. Assume G is not perfect and contains a co-triangle. By the Strong Perfect Graph Theorem, G contains an odd antihole A. Suppose no co-triangle of G contains two vertices of A. Compute the component O of $\overline{G}$ that contains A. By Lemma 3.2, in G, O is co-triangle-free and a homogeneous set, and we are done.

Now, we may assume some two vertices of A belong to a co-triangle. Define the sets Q, R, S as in Lemma 3.3. Let $G' = G[A \cup S]$. By Lemma 3.3, the graph G' is a homogeneous set of G and no co-triangle of G' contains two vertices of A. We may assume G' contains a co-triangle, for otherwise, G' is the desired homogeneous set. Now, G' satisfies the hypothesis of Lemma 3.2. Compute the component O of the complement of G' that contains A. Note that any homogeneous set of G' is a homogeneous set of G. Thus, O is a homogeneous set of G and it is co-triangle-free by Lemma 3.2. □

We note the homogeneous set in (iii) can be easily found in $O(n^3)$. Actually, with a bit more work, one can find it in $O(n^2)$ time (given the odd antihole), but this step is not the bottle neck of our recognition and coloring algorithms.

4 The 2-divisibility of (banner, odd hole)-free graphs

In this section, we prove that (banner, odd hole)-free graphs are two-divisible. Recall that a graph G is k-divisible if the vertex set of each induced subgraph H of G with at least one edge can be partitioned into k sets none of which contains a clique of size $\omega(H)$. It is easy to see that perfect graphs are 2-divisible. We will need the following two results by Hoàng and McDiarmid ([9]).

Lemma 4.1 ([9], Lemma 3) *Let G be a graph with a homogeneous set H. If for some $k \geq 2$, both H and $G[(V(G) - H) \cup \{h\}]$, for any vertex $h \in H$, are k-divisible, then so is G.* □

Theorem 4.2 ([9], Theorem 1) *A (claw, odd hole)-free graph is 2-divisible.* □

Theorem 4.3 *A (banner, odd hole)-free graph is 2-divisible.*

Proof: We prove the theorem by induction on the number of vertices. Let G be a (banner, odd hole)-free graph. By the induction hypothesis, we may assume every proper induced subgraph of G is 2-divisible. We may assume that G is not perfect, and by Theorem 4.2, G has $\alpha(G) \geq 3$. By Theorem 3.4, G contains a homogeneous set H. By the induction hypothesis, both H and $G[(V(G) - H) \cup \{h\}]$ (for any vertex $h \in H$) are 2-divisible. Now by Lemma 4.1, G is 2-divisible. □

As noted in Section 1, the statement below follows from Theorem 4.3.

Corollary 4.4 *A (banner, odd hole)-free graph G has $\chi(G) \leq 2^{\omega(G)-1}$.* □

We note the result of Scott and Seymour ([17]) showing that an odd-hole-free graph G has $\chi(G) \leq 2^{2^{\omega(G)}}$.

5 Recognition algorithms

In this section we give two polynomial-time algorithms for recognizing a (banner, odd hole)-free graph. Both algorithms rely on the following easy observation.

Observation 5.1 *Let G be a graph with a homogeneous set H. Let P be a prime induced subgraph of G. Then P cannot contain two vertices in H and some vertex in in $G - H$.* □

The following statement follows from Observation 5.1.

Observation 5.2 *Let G be a banner-free graph with a homogeneous set H. Then G contains no odd hole if and only if both H and $G[(V(G) - H) \cup \{h\}]$ contain no odd hole, for any vertex $h \in H$.* □

Let $\mathcal{B}$ denotes the class of (banner, odd hole)-free graphs. The first algorithm follows the steps suggested by Theorem 3.4. First, it checks in $O(n^5)$ time that the input graph G is (banner, C_5)-free. If $\alpha(G) \leq 2$, then G is in $\mathcal{B}$. Assume $\alpha(G) \geq 3$. Also we may assume without loss of generality that G is connected. Now, use the perfect graph recognition algorithm to find an odd hole, or odd antihole in $O(n^9)$ time. If no such hole or antihole is found, G is in $\mathcal{B}$. If an odd hole is found, G is not in $\mathcal{B}$. Assume an odd antihole A is found. Using Theorem 3.4, find a homogeneous set H containing H in $O(n^3)$ time. Recursively apply the above procedure on H and $G[(V(G) - H) \cup \{h\}]$ for a vertex $h \in H$. The input graph G is in $\mathcal{B}$ if and only if both graphs H and $G[(V(G) - H) \cup \{h\}]$ are. By Algorithm 1, the recursion is called on $O(n)$ graphs. Thus, our algorithm runs in $O(n^{10})$ time.

The second algorithm bypasses the obvious approach of the first algorithm. It uses deep properties of the modular decomposition discussed in Subsection 2.2. Observation 5.1 implies the following Lemma.

Lemma 5.3 *Let G be a banner-free graph with a modular decomposition $\mathcal{P} = (M_1, M_2, \ldots, M_\ell)$. Then G is odd-hole-free if and only if the quotient graph $G_{\mathcal{P}}$ and every M_i are odd-hole-free.* □

Given a graph G. The second algorithm starts by verifying in $O(n^5)$ time that G is (banner, C_5)-free. If $\alpha(G) \leq 2$, then G is in $\mathcal{B}$. Assume $\alpha(G) \geq 3$. Now, the algorithm finds a modular decomposition of G with the partition $\mathcal{P} = (M_1, M_2,$

$\ldots, M_\ell$). The quotient graph $G_{\mathcal{P}}$ is prime. So, by Theorem 3.4, $G_{\mathcal{P}}$ is in $\mathcal{B}$ if and only if it is perfect. Perfection of $G_{\mathcal{P}}$ can be verified in $O(n^9)$ time using the perfect graph recognition algorithm. Now, recursively verify that every M_i is odd-hole-free. Theorem 2.1 implies $\mid G_{\mathcal{P}} \mid + \sum_1^\ell \mid M_i \mid = O(n)$. Thus, the second algorithm runs in $O(n^9)$ time.

Actually, we have shown the following: Let $p(n)$ (respectively, $c(n)$, $b(n)$, $t(n)$) be the time complexity of finding an odd hole or odd antihole (respectively, C_5, banner, triangle), and let $a(n) = max(p(n), c(n), b(n), t(n))$. Then there is an $O(a(n))$-time algorithm to recognize graphs in $\mathcal{B}$. Here, we are making the reasonable assumption that $a(n)$ is at least $O(n + m)$. In the current state of knowledge, recognizing perfect graph is the bottleneck of the algorithm.

6 Optimizing (banner, odd hole)-free graphs

In this section, we consider the following four optimization problems: finding a minimum coloring, finding a minimum clique cover (the minimum coloring in the complement of the graph), finding a largest stable set, and finding a largest clique. It is interesting to note that for our class of (banner, odd hole)-free graphs, the coloring and largest stable set problems are solvable in polynomial time, whereas the clique cover and largest clique problems are NP-hard.

6.1 Finding a minimum coloring in polynomial time

In this section, we describe a robust algorithms to color a (banner, odd hole)-free graph. The perfect graph coloring algorithm of Grötschel et al. ([8]) is robust. Given a graph G, the algorithm either returns an optimal coloring of G (and a clique of the same size proving the coloring is minimum) or a correct declaration that G is not perfect.

Concerning graphs G with $\alpha(G) \leq 2$, it is well known that a minimum coloring of G can be found by finding a maximum matching M in the complement $\overline{G}$ of G. (Let t be the number of edges of M. Then we have $\chi(G) = n - t$.) The algorithm of Micali and Vazirani ([14]) finds a maximum matching of a graph in $O(m\sqrt{n})$ time.

We may assume the algorithm below (Algorithm 2: DIRECT-COLOR(G)) runs in polynomial time. DIRECT-COLOR(G) returns an optimal (minimum) coloring if G is perfect or has $\alpha(G) \leq 2$.

Let G and L be two vertex-disjoint graphs. Let G have a homogeneous set H. Define $g(G, H, L)$ to be the graph obtained from G by substituting L for H, that is, $g(G, H, L)$ is obtained from G by removing the set H, adding the graph L, and for each vertex $v \in G - H$ that has a neighbor in H, add all edges between v and L. Note that L is a homogeneous set of $g(G, H, L)$. Since the following observation is easy to establish, we omit its proof.

Algorithm 2 DIRECT-COLOR(G)

input: graph G
output: An optimal coloring of G or a message that G is not perfect.
remark: The algorithm returns an optimal coloring if G is perfect or $\alpha(G) \leq 2$.

if $\alpha(G) \leq 2$ **then**
  find an optimal coloring of G via the matching algorithm
  return the optimal coloring and stop
end if

apply the perfect graph coloring algorithm on G
if an optimal coloring is returned **then**
  return the optimal coloring and stop
else
  return "G is not perfect"
end if

Observation 6.1 *Let G be a graph with a homogeneous set H. Let $\chi(H) = k$. Let K_k be a clique on k vertices. Then*

(i) $\chi(G) = \chi(\ g(G, H, K_k)\ )$.

(ii) *if G is (banner, odd hole)-free then so is $g(G, H, K_k)$.* □

We need to discuss part (i) of Observation 6.1 before describing our algorithm. Assume we are given a minimum coloring of H with stable sets $S_1, \ldots, S_k$. Enumerate the vertices of K_k as $v_1, \ldots, v_k$. Assume in a minimum coloring of $g(G, H, K_k)$, the vertices v_i has color $c(i)$. From a minimum coloring of $g(G, H, K_k)$, we may obtain a minimum coloring of G by assigning to each vertex in S_i the color $c(i)$.

Consider a graph G with a modular decomposition $\mathcal{P} = (M_1, \ldots, M_t)$. Note that if G is prime then each M_i is a single vertex. Our coloring algorithm will return a minimum coloring of G, or correctly declare that G is not (banner, odd hole)-free. If G is prime, then a single call to DIRECT-COLOR(G) will return a minimum coloring of G, or a message that G is not (banner, odd hole)-free. Now, assume G is not prime. For each module M_i, the algorithm is recursively applied to M_i. If M_i is not (banner, odd hole)-free, then neither is G. If a minimum coloring of M_i with t colors is returned, then we substitute a clique with t vertices for M_i in G. This substitution is done for all modules of the decomposition. Let F be the graph obtained from these substitutions. By Theorem 3.4, if $\alpha(F) \geq 3$, then F cannot contain an odd antihole A, for otherwise A is contained in some maximal module of G and this module intersects some maximal module M_i of the decomposition, a contradiction. Thus a single call

to DIRECT-COLOR(F) will return a minimum coloring of F (and thus, of G) or a correct declaration that G is not (banner, odd hole)-free. Our algorithm is described below (Algorithm 3).

Algorithm 3 COLOR(G)

input: graph G with a modular decomposition $\mathcal{P} = (M_1, \ldots, M_\ell)$.
output: an optimal coloring of G or a message that G is not (banner, odd hole)-free.

if G is prime **then**
  call DIRECT-COLOR(G)
  if a minimum coloring $\mathcal{C}$ of G is returned **then**
    return $\mathcal{C}$ and stop
  else
    return "G is not (banner, odd hole)-free" and stop
  end if
end if

for each M_i of $\mathcal{P}$ **do**
  call COLOR(M_i)
  if an optimal coloring of M_i with t_i colors is returned **then**
    construct a clique C_i on t_i vertices
  else
    return "G is not (banner, odd hole)-free" and stop
  end if
end for

let F be the graph obtained from G by substituting the clique C_i for M_i for all i
call DIRECT-COLOR(F)
if an optimal coloring $\mathcal{C}$ of F is returned **then**
  from $\mathcal{C}$, construct an optimal coloring $\mathcal{C}'$ of G
  return $\mathcal{C}'$ and stop
else
  return "G is not (banner, odd hole)-free" and stop
end if

The algorithm can be make more efficient by skipping the trivial modules of the decomposition. But this does not improve the worst case complexity. Since the modules of $\mathcal{P}$ are maximal and therefore pairwise non-intersecting, the number of recursively calls to COLOR is bounded by $2n-1$. This fact can be easily seen by induction as follows: let $s_i = |M_i|$, by induction the number of recursive calls on the modules M_i is at most $(2s_1-1)+(2s_2-1)+\ldots+(2s_\ell-1) = (2\sum_{i=1}^{\ell} s_i) - \ell \leq 2n-1$. Let $g(n)$ (respectively, $t(n)$) be the time complexity of coloring a perfect graph (respectively, find a maximum matching of a graph).

By Theorem 2.1, Algorithm 3 has complexity $O(n\ (max(g(n), t(n)))$.

6.2 Finding a largest clique is hard

It was proved by Poljak ([15]) that it is NP-hard to determine $\alpha(G)$ for a triangle-free graph G. Since the proof is short, we reproduce it here. Let G be a graph. Construct a graph $f(G)$ from G by, for each edge ab of G, replacing ab by an induced path on three edges, that is, we subdivide the edge ab twice. Then it is easy to see that $\alpha(f(G)) = \alpha(G) + m$, where m is the number of edges of G. The graph $f(G)$ is triangle-free. By repeatedly applying this construction, one can show that it is NP-hard to compute $\alpha(G)$ for a graph G of any given girth g. In particular, it is NP-hard to compute $\alpha(G)$ for a (triangle, C_5)-free graph G, or equivalently, to compute $\omega(F)$ for a (co-triangle, C_5)-free graph F. Since (co-triangle, C_5)-free graphs are (banner, odd hole)-free, it is NP-hard to compute $\omega(G)$ for a (banner, odd hole)-free graph.

6.3 Finding a minimum clique cover is hard

It was noted by Jensen and Toft (section 10.3 of [11]) that it is NP-complete to decide, for any fixed integer g, whether a graph G of girth at least g is 6-colorable (this result uses the Hajós construction for graphs of high chromatic number). Thus, it is NP-hard to find a minimum coloring of a (C_3, C_4, C_5)-free graph, or equivalently, to find a minimum clique cover of a (co-triangle, $2K_2$, C_5)-free graph G. Since such a graph G is (banner, odd hole)-free, it is NP-hard to find a minimum clique cover of a (banner, odd hole)-free graph.

6.4 Finding a largest stable set in polynomial time

In this section, we will describe a polynomial-time algorithm for find a largest stable set of a (banner, odd hole)-free graph. Our algorithm is robust. We will actually solve the more general problem of finding a maximum weighted stable set.

Let G be a graph whose vertices are given "weights", that is, each vertex v is given an integer $w_G(v)$. When the context is clear, we will drop the subscript G and let $w(x) = w_G(x)$. The problem is to find a stable set S which maximizes $\sum_{v \in S} w_G(v)$, this sum is denoted by $\alpha_w(G)$. We will refer to this problem as the Maximum Weighted Stable Set problem, or MWSS for short. The algorithm of Grötschel et al. ([8]) actually solves, for a perfect graph, the "minimum weighted coloring" (which we will not define here) and the "maximum weighted clique" (which is MWSS for the complement of the graph). The algorithm of Grötschel et al. is a robust algorithm: given a weighted graph G, it either returns a maximum weighted stable set of G, or correctly declares that G is not perfect. Finally, when G is a graph with $\alpha(G) \leq 2$, then MWSS can be solved

in polynomial time since there are only $O(n^2)$ stable sets. We can list them all and take the one with the maximum weight.

Let H be a homogeneous set of a weighted graph G. By $g(G, H, h)$, we denote the weighted graph obtained from G by substituting a vertex h for H where the weight function w' for $g(G, H, h)$ is defined as follows. For the vertex h, we set $w'(h) = \alpha_w(H)$ and for all $x \in G - H$, we set $w'(x) = w(x)$. It is easy to see that $\alpha_w(G) = \alpha_{w'}(g(G, H, h))$.

Let G be a connected (banner, odd hole)-free graph with $\alpha(G) \geq 3$. If G is not perfect, then by Theorem 3.4, G contains a homogeneous set H. We recursively solve MWSS for H, then substitute a vertex h for H, where the weight of h is equal to $\alpha_w(H)$. We obtain the graph $F = g(G, H, h)$. Now, solve MWSS for F. From a maximum weighted stable set of F, we can easily construct a maximum weighted stable set of G. This discussion shows MWSS can be solved in polynomial time for (banner, odd hole)-free graphs. We can design a more efficient, and robust, algorithm by imitating the coloring algorithm in Section 6.1.

First, we construct the algorithm DIRECT-STABLE-SET (Algorithm 4) that takes as input a weighted graph G. If $\alpha(G) \leq 2$, it returns a maximum weighted stable set and terminates. Otherwise, it applies the stable set algorithm of Grötschel et al. to G. Then the algorithm either returns a maximum weighted stable set, or a message that G is not perfect.

Algorithm 4 DIRECT-STABLE-SET(G)

input: a graph G with a weight function w on its vertices.
output: a maximum weighted stable set of G or a message that G is not perfect.
remark: The algorithm returns a maximum stable set of G whenever G is perfect or $\alpha(G) \leq 2$.

if $\alpha(G) \leq 2$ **then**
  find maximum weighted stable set S of G by listing all stable sets of G and taking one with the largest weight
  return S and stop
end if

apply the stable set algorithm for perfect graph on G
if a maximum weighted stable set S is returned **then**
  return S and stop
else
  return "G is not perfect"
end if

Consider a graph G with a modular decomposition $\mathcal{P} = (M_1, \ldots)$. The algorithm returns a maximum weighted stable set of G, or correctly declares

that G is not (banner, odd hole)-free. If G is prime, then a single call to DIRECT-STABLE-SET(G) either returns a maximum weighted stable set of G, or a message that G is not (banner, odd hole)-free. Now, consider the case G is not prime. For each module M_i, the algorithm is recursively applied on M_i. If M_i is not (banner, odd hole)-free, then neither is G. If a maximum weighted stable set S_i is returned, then the algorithms substitutes a vertex h_i for M_i in G and gives h_i a weight equal to $\alpha_w(M_i)$. This substitution is done for all modules of the decomposition. Let F be the graph obtained from these substitutions. Note that F is the quotient graph $G_{\mathcal{P}}$ and so it is prime. By Theorem 3.4, if $\alpha(F) \geq 3$, then F cannot contain an odd antihole A. Thus a single call to DIRECT-STABLE-SET(F) will return a maximum weighted stable set of F (from this we can construct a maximum weighted stable set of G in the obvious way) or a correct declaration that G is not (banner, odd hole)-free. Our algorithm is described below (Algorithm 5). Let $p(n)$ be the time complexity of finding a maximum weighted stable set of a perfect graph. Our algorithm runs in $O(p(n) + n^2)$ time with the term n^2 coming from solving MWSS for graph with no co-triangles.

References

[1] C. Berge. Les problèmes de colorations en théorie des graphes. *Publ. Inst. Statist. Univ. Paris IX*, pages 123–160, 1960.

[2] M. Chudnovsky, G. Cornuéjols, X. Liu, P. Seymour, and K. Vuskovic. Recognizing berge graphs. *Combinatorica*, 25:143–186, 2005.

[3] M. Chudnovsky, N. Robertson, P. Seymour, and R. Thomas. The strong perfect graph theorem. *Annals of Mathematics*, 64:51–229, 2006.

[4] V. Chvátal and N. Sbihi. Recognizing claw-free perfect graphs. *Journal of Combinatorial Theory, Series B*, 44(2):154–176, 1988.

[5] M. Conforti, G. Cornuéjols, A. Kapoor, and K. Vušković. Decomposition of odd-hole-free graphs by double star cutsets and 2-joins. *Discrete Applied Mathematics*, 141:41–91, 2004.

[6] T. Gallai. Transitiv orientierbare graphen. *Acta Math. Acad. Sci. Hungar*, 18:25–66, 1967.

[7] M. U. Gerber, A. Hertz, and V. V. Lozin. Stable sets in two subclasses of banner-free graphs. *Discrete Applied Mathematics*, 132:121–136, 2004.

[8] M. Grötschel, L. Lovász, and A. Schrijver. Polynomial algorithms for perfect graphs. In Berge and Chvátal, editors, *Topics on perfect graphs*, pages 325–356. North-Holland Mathematics Studies, 1984.

[9] C. T. Hoàng and C. McDiarmid. On the divisibility of graphs. *Discrete Mathematics*, 242:145–156, 2002.

[10] C. T. Hoàng and B. A. Reed. Some classes of perfectly orderable graphs. *Journal of Graph Theory*, 13(4):445–463, 1989.

[11] T. R. Jensen and B. Toft. *Graph coloring problems*. Wiley-Interscience, 1995.

[12] L. Lovász. Normal hypergraphs and the perfect graph conjecture. *Discrete Math*, 2:253–267, 1972.

[13] R. M. McConnell and J. Spinrad. Linear-time modular decomposition and efficient transitive orientation of comparability graphs. In *SODA*, volume 94, pages 536–545, 1994.

[14] S. Micali and V. V. Vazirani. An $O\left(\sqrt{|V|}|E|\right)$ algorithm for finding maximum matching in general graphs. In *Foundations of Computer Science, 1980., 21st Annual Symposium on*, pages 17–27. IEEE, 1980.

[15] S. Poljak. A note on stable sets and colorings of graphs. *Commentationes Mathematicae Universitatis Carolinae*, 15(2):307–309, 1974.

[16] V. Raghavan and J. Spinrad. Robust algorithms for restricted domains. *Journal of Algorithms*, 48:160–172, 2003.

[17] A. Scott and P. Seymour. Induced subgraphs of graphs with large chromatic number. I. Odd holes. http://arxiv.org/abs/1410.4118, 2015.

[18] Z. Tuza, D. Král, J. Kratochvíl, and G. J. Woeginger. Complexity of coloring graphs without forbidden induced subgraphs. In *WG 2001, LNCS 2204*, pages 254–262, 2001.

Algorithm 5 MWSS(G)

input: graph G with a modular decomposition $\mathcal{P} = (M_1, \ldots, M_\ell)$ and a weight function w on its vertices.
output: a maximum weighted stable set of G or a message that G is not (banner, odd hole)-free.

if G is prime **then**
  call DIRECT-STABLE-SET(G)
  if a maximum weighted stable set S of G is returned **then**
    return S and stop
  else
    return "G is not (banner, odd hole)-free" and stop
  end if
end if

for each M_i of $\mathcal{P}$ **do**
  call MWSS(M_i)
  if a maximum weighted stable set S_i of M_i is returned **then**
    compute $\alpha_w(M_i)$ which is the sum of the weights of the vertices in S_i
  else
    return "G is not (banner, odd hole)-free" and stop
  end if
end for

Let F be the graph obtained from G by substituting a vertex s_i for M_i for all i, and set $w(s_i) = \alpha_w(M_i)$
call DIRECT-STABLE-SET(F)
if a maximum weighted stable set S of F is returned **then**
  construct a maximum weighted stable set $S*$ of G from S
  return $S*$ and stop
else
  return "G is not (banner, odd hole)-free" and stop
end if